The Analysis of Publish/Subscribe Systems over Mobile Wireless Ad Hoc Networks*

Thadpong Pongthawornkamol, Klara Nahrstedt
University of Illinois at Urbana-Champaign
{tpongth2,klara}@cs.uiuc.edu

Guijun Wang
Boeing Phantom Works
guijun.wang@boeing.com

Technical Report to Boeing Phantom Works. The article was published in Mobiquitous 2007.

Abstract

In this work, we analyze the publish/subscribe distributed system paradigm over mobile ad hoc networks with respect to the performance and the impact of different mobility speeds and patterns of all participants in urban and high-way scenarios. In order to handle mobility, our publish/subscribe framework relies on light-weight publish/subscribe brokers, which are chosen in a topic-driven fashion. With the option of geographic location and waypoint information, we discuss several publish/subscribe variants and analyze their advantages/disadvantages when applied in different mobility environment. Via simulations, we measure quality of service in terms of delivery ratio and message delay achieved by such different publish/subscribe schemes over different scenarios and mobility models.

1 Introduction

Publish/subscribe architecture is a group communication paradigm where information consumers (i.e. *subscribers*) and producers (i.e. *publishers*) with mutual interests are linked altogether by publish/subscribe intermediate brokers. In publish/subscribe systems, subscribers specify characteristics of their interests (i.e. , *topics*) to publish/subscribe brokers. Upon receiving information from publishers, the publish/subscribe brokers then filter and deliver information to subscribers whose interests match the topic and content of the information. Publish/subscribe architecture offers several scalable properties including time-, space-, and synchronization- decoupling [1]. It allows information to be disseminated from publishers to subscribers asynchronously without having subscribers or publishers aware of other communication endpoints. Such properties allow publish/subscribe systems to become communication backbones in large-scale, dynamic networks.

While most of previous work on publish/subscribe systems have been focusing on ensuring reliability and improving efficiency in publish/subscribe systems over traditional wired, Internet-based network architecture, there are few works on publish/subscribe systems over wireless, especially wireless ad-hoc networks. Wireless ad-hoc networks possess several unique characteristics that could represent profit toward or hindrance against publish/subscribe systems. For example, broadcast medium in wireless networks allow a single message to reach multiple recipients with only one transmission. At the same time, high node mobility in wireless ad hoc networks usually causes network partition or out-of-date routing information. Current approaches in publish/subscribe systems assume stability of underlying physical network topology, which does not hold in wireless ad hoc networks. For example, one might want to build a publish/subscribe system over a vehicular ad hoc networks (VAN) to distribute traffic information on a highway or local city information in metropolitan area. Another example is a post-disaster recovery operation where rescue officers and their vehicles move around the site and share any event they have found.

In order to achieve efficient, reliable and flexible group communication and data dissemination over wireless ad hoc networks, two dynamism problems must be addressed: *logical mobility* and *physical mobility* [2]. Logical mobility is

*This project is funded by Boeing Phantom Works.

application-level mobility caused by communication endpoints that change their information interests. Logical mobility problem can be addressed with the traditional publish/subscribe architecture. On the other hand, the traditional publish/subscribe paradigm does not address the problem of physical mobility, which is caused by both communication endpoints and publish/subscribe intermediate brokers to change their physical locations, resulting in topology changes and route breaks. The severity of the physical mobility problem depends on the movement speed of nodes in the networks. To our knowledge, such problem has not been well addressed by current publish/subscribe technologies. Thus, one might come up with a simple but non-trivial existential question : *Is it feasible to implement a reliable publish/subscribe system over wireless ad hoc networks?*

In order to answer such a question, several sub-problems remain to be answered. Is current publish/subscribe technology, which is designed to address logical mobility, suitable for physically-dynamic wireless ad-hoc networks? What can be potential problems in setting up a traditional publish/subscribe system over wireless networks? Can we improve the publish/subscribe system's performance by the use of GPS-equipped nodes in wireless networks? To which extent of mobility can a publish/subscribe system be sustained? Answering these questions would give us a better understanding of how to develop a reliable publish/subscribe system over wireless ad-hoc networks.

In this paper, we analyze the publish/subscribe distributed system paradigm over mobile ad hoc networks with respect to the performance and the impact of different node mobility speeds in the pub/sub system in urban and high-ways scenarios. Our work has three contributions. First, we discuss the potential reliability problems in publish/subscribe systems over wireless ad hoc networks. Second, we discuss several design alternatives of publish/subscribe frameworks to solve the mobility problem. Finally, we perform extensive evaluation via simulations to investigate the feasibility of creating publish/subscribe systems over several scenarios of wireless ad hoc networks.

The rest of this paper is organized as follows. Section 2 presents model and assumptions used in our work. Section 3 describes the basic ideas of our design alternatives of publish/subscribe systems. Section 4 illustrates our publish/subscribe systems in several scenarios. Section 5 shows the experimental results. Section 6 discusses some related works on publish/subscribe systems over wireless ad hoc networks. Finally, Section 7 concludes the paper.

2 System Model and Assumptions

Our work assumes the following model and assumptions.

Network Model: We model the system as a wireless ad hoc network consisting of only mobile nodes that are free to move independently. We do not assume any fixed infrastructure. Any data will be routed from its source node through intermediate nodes to the destination node using wireless ad hoc routing protocols. In this paper, we assume all communications to be between nodes within the ad hoc network. However, simple modifications can be made to combine our system with infrastructure-based networks.

Node Model: Each node is equipped with a GPS device, which enables each node to obtain its current geographical location. Each node has its own unique identifier. In this paper, we focus on performance aspects of publish/subscribe systems such as delivery ratio and delay. Hence, we do not focus on security aspects such as confidentiality and integrity of messages. Several existing techniques such as node certificates and key encryptions can be used to address such issues and are beyond the scope of this paper.

Publish/Subscribe Model: Each node can be either a *subscriber*, a *publisher*, or both. We do not assume static publish/subscribe broker infrastructure. We also assume content-based publish/subscribe model, where a policy is associated with each published event, and each client embeds a *topic* of interest and a filter predicate along with its subscription. A published event will be delivered to a subscriber if and only if *all* of these conditions are true : 1) the topic of the event matches the topic that subscriber specified in its subscription, 2) the content of the event satisfies the filter predicate specified in the subscriber's subscription, and 3) the subscriber itself has attributes which satisfies the access control policy embedded with that event.

3 Design Alternatives in Publish/Subscribe System

In this section, we discuss main ideas and basic operations of our several design alternatives in our publish/subscribe system. The main focus of our work is to increase system's quality of service in terms of message reliability under the

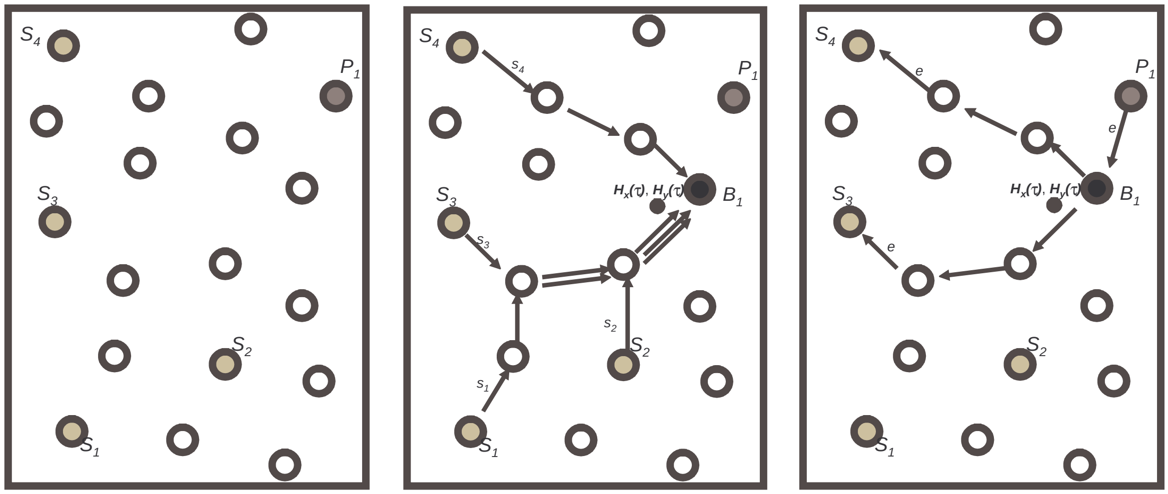

(a) Four subscribers and one publisher are interested in topic τ

(b) Subscribers send subscriptions to the broker node

(c) The publisher publishes an event to the broker node, which then evaluates policy and deliver event to eligible recipients.

Figure 1. Basic operations of the broker-mode geographic-based publish/subscribe system. There are 4 subscribers (S_1, S_2, S_3, S_4) and 1 publisher (P_1) whose topic of interest is τ. Only S_3 and S_4 have predicates that match the content of event e. All point-to-point communications are carried on by geographic-based ad hoc routing protocol.

presence of node mobility. We present analysis on several design choices, which will be validated in Section 5

3.1 Routing Protocols

A publish/subscribe protocol can rely on different routing protocols. We will investigate two different routing protocols, GPSR [3] and AODV [4], for the publish/subscribe systems. In geographic-based routing protocol, the destination node of any message is specified by its node identifier and its *geographical coordinate*. Using periodic beaconing scheme, each node maintains geographic locations of itself and its one-hop neighbors. In normal operations, the routing is performed using *greedy forwarding* as follows. Upon receiving any packet, a receiving node picks up a neighbor whose location is closest to the final location specified in the packet, and forwards the packet to that neighbor. In some circumstances, however, the receiving node itself is the node closest to the destination, and thus the receiving node cannot perform greedy forwarding. In such situations, special heuristics must be used to route the packet away from the local maxima towards the final destination. For example, when the greedy forwarding mode is not applicable, the forwarding node in GPSR protocol will switch its operation to *perimeter* mode, which will route the packet counter-clockwise around the void (the area in which the greedy forwarding mode is not effective) until the packet escapes the local maxima or the packet return to the same forwarding node. In the latter case, that packet will be dropped, since the destination node does not locate at the position specified in the packet.

A publish/subscribe system can leverage the use of geographic-based routing protocols, specifically GPSR, to route all application-level messages (i.e. advertisements, subscriptions, published events) to their intended receivers. Compared to traditional node-based wireless ad hoc routing protocols, geographic-based routing protocols are more resilient to node mobility as the routing path of each message is not fixed as in node-based wireless ad hoc routing protocols such as AODV. However, AODV is a reactive, on-demand, node-based routing protocol. Whenever a node wants to send a message, it floods the route request to find the path to the destination. The flooding overhead is reduced by route caching.

3.2 Peer-based Policy Decision Point

Since we do not assume any dedicated publish/subscribe broker infrastructure due to high overhead in maintenance, the responsibility for evaluating policy of each message must be placed at mobile nodes in a peer-to-peer manner. There are two

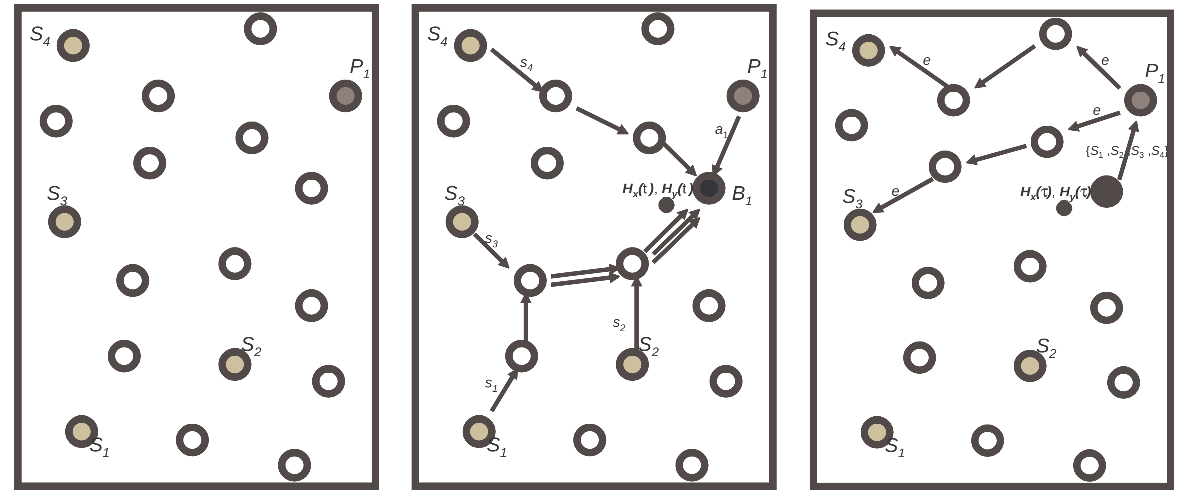

(a) Four subscribers and one publisher are interested in topic τ

(b) Subscribers send subscriptions to the rendezvous node. The publisher sends the advertisement to rendezvous node, which then sends back the subscription list back to the publisher

(c) The publisher evaluates policy with the obtained subscription list and sends the event directly to the subscribers

Figure 2. Basic operations of the rendezvous-mode geographic-based publish/subscribe system. There are 4 subscribers (S_1, S_2, S_3, S_4) and 1 publisher (P_1) whose topic of interest is τ. Only S_3 and S_4 have predicates that match the content of event e. All point-to-point communications are carried on by geographic-based ad hoc routing protocol.

viable alternatives we discuss in this paper. The first approach is called the *broker* approach and the second approach is called the *rendezvous* approach. The two approaches are different in where the policy evaluation is performed.

3.2.1 Broker-based Publish/Subscribe System

In the broker-based publish/subscribe system, each publisher will relay a published event to subscribers whose interest matches its publishing topic through one or more *broker* nodes. A broker node is a mobile node which matches / forwards events obtained from publishers to appropriate subscribers. A broker node can be any arbitrary node. However, to ensure validity of the publish/subscribe protocol, publishers and subscribers who share the same topic must select at least one common broker. The process to select the broker for each topic will be discussed in Section 3.3.

In the broker-based publish/subscribe system, each subscriber periodically sends a subscription message containing its identifier, its topic of interest, and its geographical location. Upon receiving a subscription, a broker stores the subscription in its database. Whenever a publisher publishes an event, it sends the published event to its chosen broker nodes. The broker nodes then perform policy evaluation with the each subscription in the broker's database, one by one. If a subscriber's attributes satisfy the policy, the broker nodes send the message directly to the subscriber using the geographic-based routing protocol.

3.2.2 Rendezvous-based Publish/Subscribe System

In contrast to the broker-based publish/subscribe system, the policy evaluation in the rendezvous-based publish/subscribe system is done at the publisher side. When a publisher wants to publish an event, it performs the policy evaluation to the list of subscriptions obtained from *rendezvous* nodes and sends the event directly to the subscriber. To get the list of interested subscribers, a publisher periodically sends *advertisement* message to selected rendezvous nodes. Upon receiving an advertisement from a publisher, a rendezvous node performs topic matching (without policy evaluation) between the advertisement and the set of subscriptions that rendezvous node has received, and sends the list of subscriptions with the same topic back to the advertising publisher. Thus, the rendezvous node acts as a directory service for each publisher to retrieve the list of subscribers that share the common topic of interest. Similar to the broker-based publish/subscribe system, publishers and subscribers who share the same topic must have at least one common rendezvous node.

Since each publisher in the rendezvous-based publish/subscribe system can send the event message directly to each subscriber without the need for forwarding brokers, the rendezvous-based publish/subscribe system tends to consume less bandwidth resource and have less delivery latency than the broker-based publish/subscribe system. However, the rendezvous-based publish/subscribe system is more sensitive to node mobility since each subscription takes longer time to arrive at the policy evaluation node.

3.3 Broker/Rendezvous Node Selection

To ensure validity of the publish/subscribe system, publishers and subscribers with the same topic must share at least one common broker / rendezvous node. One simple selection technique to achieve such goal is *node-based technique*, which uses a globally defined *hash function* to hash the topic into node identifier and use *node whose identifier is closest to the hashed value* to become the broker / rendezvous node of that topic. However, a location of the node whose identifier closest to the hashed value must be obtained using a separate location lookup system to route the subscriptions and advertisements to that node.

Another approach is to use the *geographic-based technique* of the hash function scheme as follows. Every node possesses two globally defined hash functions, H_x and H_y, which use topic as an input and return geographic coordinate in x-axis and y-axis respectively. A subscriber or publisher with topic of interest τ will periodically send subscriptions or advertisements to location $(H_x(\tau), H_y(\tau))$. A node whose current location closest to $(H_x(\tau), H_y(\tau))$ will act as a broker / rendezvous node of topic τ. This approach utilizes the benefit of geographic-based routing protocol without the need for separate location lookup system. The following subsection will show the example of operations in our publish/subscribe system.

3.4 Basic Operations

To illustrate the basic operations of the publish/subscribe system with geographic-based variant, Figure 1 shows an example of the broker-based publish/subscribe system consisting of a publisher (P_1) and four subscribers (S_1,S_2,S_3,S_4) sharing a common topic τ. Each subscriber periodically sends the subscription specifying the topic of interest and the filtering predicate to the location $(H_x(\tau), H_y(\tau))$. Using GPSR routing protocol, the subscription will be routed to the node B_1 whose location is closest to $(H_x(\tau), H_y(\tau))$. B_1 then becomes a broker for topic τ. Once receiving subscriptions s_1, s_2, s_3, s_4, B_1 stores those subscriptions in its subscription table. Subsequently, publisher P_1 publishes an event e with topic τ. The event e is then routed to the same location $(H_x(\tau), H_y(\tau))$ and captured by B_1. B_1, which then performs the policy evaluation and delivers e to interested subscribers. We assume that only S_3 and S_4 are subscribers whose predicates match the content in e

As well as the broker-based publish/subscribe system, Figure 2 shows the same scenario in the rendezvous-based publish/subscribe system. In addition to periodic subscriptions sent by each subscriber, the publisher P_1 also periodically sends the advertisement a to keep itself updated with the current subscription list with topic τ. When P_1 publishes an event e, it can perform the policy evaluation and send the event directly to S_3 and S_4.

It should be noted that the operations of node-based publish/subscribe systems are similar to geographic-based variants except that the topic of interest is hashed into node identifier instead. A node-based routing protocol (such as DSR or AODV) is then used as an underlying point-to-point communication mechanism. Hence, unless a failure occurs, the broker / rendezvous node of a topic is fixed over time in node-based publish/subscribe systems. On the other hand, the broker / rendezvous node of a topic can change over time due to node mobility.

Any logical mobility in the system will be handled by soft-state information handling. Using periodic subscriptions, a new broker / rendezvous node will receive the subscriptions from interested subscriber. Any subscriber who fails to send the subscription to the broker/rendezvous node will be removed from the broker/rendezvous node's subscription list using time-out mechanism. With this approach, there is no need to transfer any state from the old rendezvous node to the new one. Also any failures occurred to subscribers or rendezvous nodes will not disrupt the publish/subscribe service.

3.5 Proximity Broker / Rendezvous Node Selection in Geographic-based Publish/Subscribe Systems

In the basic approach of the geographic-based publish/subscribe system, the node whose geographic location closest to the hashed coordinate values of a topic will become the broker / rendezvous node of that topic. In the scenarios where node density and node mobility is high, such basic approach could lead to frequent broker / rendezvous changes or even

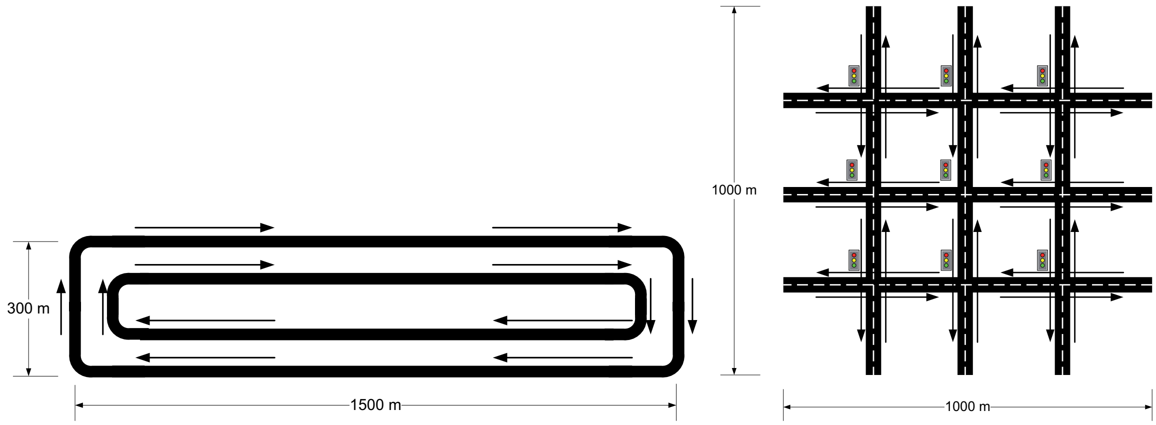

(a) Highway Traffic Event Reporting (b) Manhattan mobility model

Figure 3. Mobility scenario maps

oscillations, which can cause each event to be delivered to only partial set of interested subscribers. To address such mobility issue, we use the concept of *proximity broker / rendezvous node selection* that allows any node in the proximity of the hashed coordinate values to become broker / rendezvous node. Care must be taken, however, to prevent duplicate event delivery at subscribers. The proximity broker / rendezvous node selection is done as follows.

In a broker-based publish/subscribe system, when a node U receives a subscription s with topic τ and hashed destination location $(H_x(\tau), H_y(\tau))$, and U's geographical location is within globally predefined distance K from coordinate $(H_x(\tau), H_y(\tau))$, U then captures the subscription s and broadcast s to all of its neighbors. Upon receiving the broadcast of s, each of U's neighbors also stores s in its database. When a node V receives an event e with topic τ and finds that its location is within K from e's hashed destination, V captures e (without broadcasting) and processes e with its subscription database. By setting K to an appropriate value (i.e. $K =$ radio transmission range divided by 2), all nodes within range K from coordinate $(H_x(\tau), H_y(\tau))$ will receive all subscriptions with topic τ. Moreover, only one of such nodes will process each event with topic τ with the complete subscription list.

As well as the broker-based publish/subscribe system, the same procedure applies to a rendezvous-based publish/subscribe system, but with proximity advertisement processing instead of proximity event processing. Each event will still be processed at its source publisher, but a publisher will receive a subscription list from any but only one node within range K from the hashed coordinate.

3.6 Geographic-based Publish/Subscribe Systems with Waypoint Information

As described in Section 3.3, the basic concept of geographic-based broker node selection is to hash the topic of interest into a two-dimension coordinate. However, such a basic scheme can suffer from scenarios where geographic distribution of nodes is not uniform. For example, nodes that participate in a publish/subscribe system in a highway will always follow the highway path only. Thus, hashing the topic uniformly will be likely to return a geographic location that is not in the highway track, causing the geographic routing protocol to route messages to a void location. In high-mobility scenarios, such scheme could lead to significant performance degradation.

With additional waypoint information (i.e. road maps), a better geographic-based broker node selection scheme can be done by hashing the topic only into a coordinate in the waypoint track. One simple way to do so is to use the basic uniform hash function to hash the topic into a coordinate, and then choose the waypoint that is closest to the hashed coordinate as the home zone of that topic. Hashing to waypoint can reduce the possibility to map a topic into void areas, increasing reliability of event delivery.

4 Scenarios

We envision the use of our approaches in several application scenarios. Here, we present two of such scenarios and how the geographic-based and node-based publish/subscribe systems can be applied to them.

Parameters	Highway Traffic Report	Police Surveillance
#Nodes	60	60
Mobility model	Freeway (Figure 3(a))	Manhattan (Figure 3(b))
Speed (m / s)	10 - 60	1 - 20
Message size (bytes)	512	1024
Publish interval (secs)	10	10
#Topics	2	6
#Publishers per topic	1	2
#Subscribers per topic	9	5

Table 1. Simulation Parameters

Scenario 1 : Highway Traffic Reporting : A geographic-based publish/subscribe system can be used to report traffic condition in a highway. In this scenario, each publisher is a vehicle which periodically measures its speed and location. Such information can be sent to subscribers, which are other vehicles in order to allow them to know the traffic condition at the publisher's location. Different subscribers may be interested in different event conditions. For example, a subscriber vehicle may be interested to receive events indicating that the speed of the vehicle in the another highway is more than a threshold (e.g. a better path to switch) while another subscriber is interested in events indicating that the speed of the vehicle ahead in the same direction has dropped to zero (e.g. an accident ahead).

It can be seen that the highway traffic reporting scenario depicts the freeway mobility model. Each node in the highway moves with high speed but with the same direction. Figure 3(a) depicts one example of freeway maps.

Scenario 2 : Police Surveillance in Metropolitan Area : Another scenario to which the publish/subscribe system can be applied is police surveillance in metropolitan area. In such scenario, a publisher is a cooperative vehicle that reports any accident or criminal activity that vehicle encounters. A subscriber is a police car patrolling around the city. In this case, the publisher reports its location, time, type and severity of the event it encounters. Each subscriber may have different interests in receiving events of different types or different severity.

The police surveillance scenario then can be modeled as a Manhattan network. Each node's mobility is relatively lower than the freeway model, since each vehicle's speed is limited in a city. Figure 3(b) shows an example of Manhattan network maps.

5 Evaluation Results

We evaluate our publish/subscribe systems via simulations under two scenarios described in Section 4 using NS-2 network simulator tool [5]. We compare performance between geographic-based approach on top of GPSR [3] and the node-based publish/subscribe system on top of AODV [4] ad hoc routing protocol. Furthermore, the geographic-based approach is divided into basic variant and waypoint variant (described in section 3.6). Hence, there are totally six publish/subscribe schemes tested in the simulation : rendezvous-node-based variant (denoted by R-AODV), rendezvous-basic-geographic-based variant (denoted by R-GPSR(no map)), rendezvous-waypoint-geographic-based variant (denoted by R-GPSR(with map)), broker-node-based variant (denoted by B-AODV), broker-basic-geographic-based variant (denoted by B-GPSR(no map)), and broker-waypoint-geographic-based variant (denoted by B-GPSR(with map)). In each simulation, each subscriber node sends a subscription message every second. Each simulation run for 900 seconds. The performance is then measured in terms of delivery reliability, end-to-end delay, and routing-level message load. Each result presented in this paper is the averages from three simulations with three different seeds.

Event Delivery Reliability : Figure 4(a) and Figure 5(a) shows the delivery reliability of the six schemes in freeway mobility scenario and Manhattan mobility scenario respectively. As seen from the figures, geographic-based variants without the waypoint extension yield very low reliability in both scenarios since most topics are hashed to void regions. However, adding the waypoint extension increase reliability in geographic-based variants significantly because hashing topics into waypoint locations ensures no void destination. In most settings, broker-waypoint-geographic-based variant gives the best reliability among all variants. In addition, the rendezvous-geographic-based variants are more sensitive to mobility because the subscribers' location information gets stale more quickly.

End-to-end Event Delivery Delay: Figure 4(b) and Figure 5(b) show the average end-to-end delivery delay of the six schemes in freeway mobility scenario and Manhattan mobility scenario respectively. The graph indicates the stability problem in GPSR routing protocol without waypoint information, since packets spend most of the time to escape the void. On the

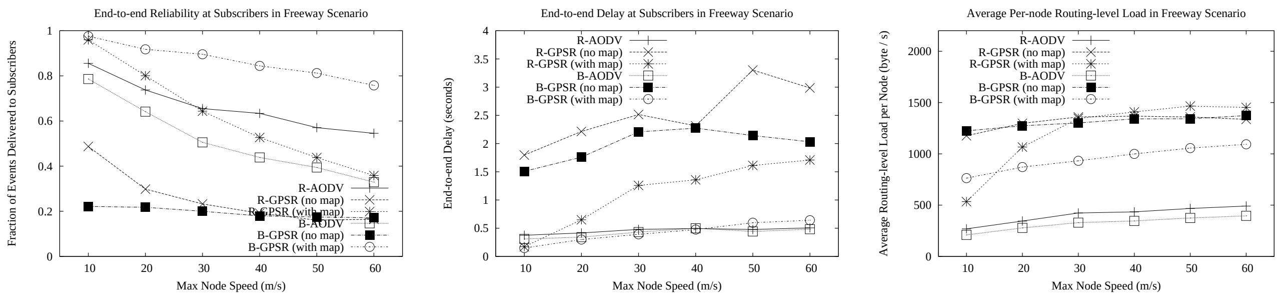

(a) Message reliability (b) End-to-end delivery delay (c) Average per-node routing-level load

Figure 4. Freeway mobility scenario

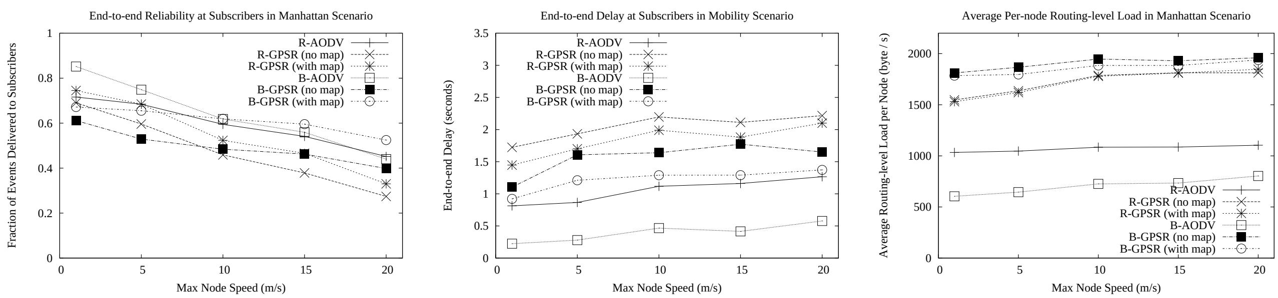

(a) Message reliability (b) End-to-end delivery delay (c) Average per-node routing-level load

Figure 5. Manhattan mobility scenario

other hand, the AODV variants give relatively lower delay even when the mobility is high. The rendezvous-geographic-based schemes are most sensitive to mobility because it takes relatively longer time to send the subscribers' locations to the publisher before an event can be sent from the publisher. However, unlike other geographic-based schemes, the broker-waypoint-geographic-based scheme give relatively smaller delay.

Per-node Message Load: Figure 4(c) and Figure 5(c) show the average per-node routing-level bandwidth consumption of the six schemes in freeway mobility scenario and Manhattan mobility scenario respectively. It can be seen that, by overall, the bandwidth consumed by proactive geographic-based schemes are much more that one consumed by reactive node-based schemes. This is because periodic beacon messages used in proactive scheme. However, we expect the difference to be less significant as the system load increases. In both scenarios, the broker-based scheme consumes less bandwidth since no extra topic advertisements needs to be sent from publishers to brokers.

6 Related Work

Most of the previous works on publish/subscribe systems [6–9] have been focusing on increasing reliability and efficiency in publish/subscribe systems over traditional fixed infrastructure networks (i.e. Internet). Such systems do not consider physical mobility issue, which arises in wireless ad hoc networks. [10] discusses techniques to map a logical publish/subscribe broker tree in physical wireless ad hoc network. However, [10] assume the ad hoc network to have low mobility and long pause time. [2] propose a publish/subscribe framework that address the mobility of client with the assumption of fixed broker infrastructure. [11] presents a publish/subscribe system on top of completely ad hoc networks but they assume subscribers to be interested in topics in their close proximity.

With the use of global positioning systems (GPS) and geographic-based routing protocols such as GPSR [3], a new paradigm of binding information to a geographical location instead of a mobile node has spawned several applications for wireless ad hoc networks. For example, [12] propose a geographic distributed hash table on top of GPSR routing protocol. Our geographic-based publish/subscribe system relies on the same approach, but in the context of publish/subscribe systems. Another use of geographic location to improve quality of service in publish/subscribe systems can be found in [13]. In contrast to [13] that divides area into cells and route messages along adjacent cell, our work uses GPSR to route messages directly, which is guaranteed to be able to route any messages if a path between sender and receiver exists.

7 Conclusion

In this paper, we discussed the mapping of publish/subscribe paradigm into wireless ad hoc networks and its potential problem. We justified several design alternatives (geographic-based V.S. node-based routing protocol, rendezvous-based V.S. broker-based event matching, with V.S. without waypoint information) along with their advantages/disadvantages. Finally, we performed comparative evaluations of such several publish/subscribe schemes in Freeway and Manhattan mobility model. The results indicated that the node-based publish/subscribe schemes perform generally better than their geographic-based counterparts. However, incorporating waypoint information into geographic-based schemes significantly improves performance. In conclusion, the publish/subscribe paradigm can be applied to wireless ad-hoc network under various mobility models. Moreover, additional location and waypoint information can further improve the performance of the system, regardless of mobility speed and pattern.

References

[1] P. T. Eugster, P. Felber, R. Guerraoui, and A.-M. Kermarrec, "The Many Faces of Publish/Subscribe." *ACM Comput. Surv.*, vol. 35, no. 2, pp. 114–131, 2003.

[2] L. Fiege, F. C. Gärtner, O. Kasten, and A. Zeidler, "Supporting Mobility in Content-Based Publish/Subscribe Middleware." in *Proc. ACM Middleware '03*, ser. Lecture Notes in Computer Science, M. Endler and D. C. Schmidt, Eds., vol. 2672. Springer, 2003, pp. 103–122.

[3] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks." in *Proc. ACM MOBICOM' 00*, 2000, pp. 243–254.

[4] C. E. Perkins and E. M. Belding-Royer, "Ad-hoc On-Demand Distance Vector Routing." in *WMCSA*. IEEE Computer Society, 1999, pp. 90–100.

[5] "NS-2 Network Simulator," http://www.isi.edu/nsnam/ns/.

[6] P. Pietzuch and J. Bacon, "Hermes: A Distributed Event-Based Middleware Architecture." in *Proc. DEBS'02*, 2002.

[7] J. Bacon, K. Moody, J. Bates, R. Hayton, C. Ma, A. McNeil, O. Seidel, and M. D. Spiteri, "Generic support for distributed applications." *IEEE Computer*, vol. 33, no. 3, pp. 68–76, 2000.

[8] A. Carzaniga, D. S. Rosenblum, and A. L. Wolf, "Design and evaluation of a wide-area event notification service." *ACM Trans. Comput. Syst.*, vol. 19, no. 3, pp. 332–383, 2001.

[9] G. Cugola, E. D. Nitto, and A. Fuggetta, "The JEDI Event-Based Infrastructure and Its Application to the Development of the OPSS WFMS." *IEEE Trans. Software Eng.*, vol. 27, no. 9, pp. 827–850, 2001.

[10] Y. Huang and H. Garcia-Molina, "Publish/Subscribe Tree Construction in Wireless Ad-Hoc Networks." in *Mobile Data Management*, ser. Lecture Notes in Computer Science, M.-S. Chen, P. K. Chrysanthis, M. Sloman, and A. B. Zaslavsky, Eds., vol. 2574. Springer, 2003, pp. 122–140.

[11] R. Meier and V. Cahill, "STEAM: Event-based Middleware for Wireless Ad Hoc Network." in *ICDCS Workshops*. IEEE Computer Society, 2002, pp. 639–644.

[12] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, and S. Shenker, "GHT: a Geographic Hash Table for Data-centric Storage." in *WSNA*, C. S. Raghavendra and K. M. Sivalingam, Eds. ACM, 2002, pp. 78–87.

[13] N. Carvalho, F. Araújo, and L. Rodrigues, "Reducing Latency in Rendezvous-Based Publish-Subscribe Systems for Wireless Ad Hoc Networks." in *ICDCS Workshops*. IEEE Computer Society, 2006, p. 28.